Fairness Beyond Disparate Treatment & Disparate Impact: Learning Classification without Disparate Mistreatment

Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, Krishna P. Gummadi
MPI-SWS, Germany
{mzafar, ivalera, manuelgr, gummadi}@mpi-sws.org

arXiv:1610.08452v1 [stat.ML] 26 Oct 2016

ABSTRACT

As the use of automated decision making systems becomes wide-spread, there is a growing concern about their potential unfairness towards people with certain traits. Anti-discrimination laws in various countries prohibit unfair treatment of individuals based on specific traits, also called sensitive attributes (*e.g.*, gender, race). In many learning scenarios, the trained algorithms (classifiers) make decisions with certain inaccuracy (misclassification rate). As learning mechanisms target minimizing the error rate for *all* decisions, it is quite possible that the optimally trained algorithm makes decisions for users belonging to different sensitive attribute groups with different error rates (e.g., decision errors for females are higher than for males). To account for and avoid such unfairness when learning, in this paper, we introduce a new notion of unfairness, *disparate mistreatment*, which is defined in terms of misclassification rates. We then propose an intuitive measure of disparate mistreatment for decision boundary-based classifiers, which can be easily incorporated into their formulation as a convex-concave constraint. Experiments on synthetic as well as real world datasets show that our methodology is effective at avoiding disparate mistreatment, often at a small cost in terms of accuracy.

1. INTRODUCTION

The emergence and wide-spread usage of automated data-driven decision making systems in a wide variety of applications, ranging from content recommendations to pretrial risk assessment, has raised concerns about their potential unfairness towards people with certain traits [8, 11, 17]. Anti-discrimination laws in various countries prohibit unfair treatment of individuals based on specific traits, also called *sensitive* attributes (*e.g.*, gender, race). These laws typically distinguish between two different notions of unfairness [4]: ***disparate treatment***, which focuses on the decision making *process* itself, and ***disparate impact***, which focuses on the *outcomes* of the process. More specifically, there is disparate treatment if the decision making process uses sensitive attribute information, and there is disparate impact if the outcomes of the decision making process disproportionately benefit or hurt members of certain sensitive attribute value groups. A number of recent studies have focused on designing decision making systems which avoid one or both of the above types of unfairness [9, 12, 13, 22].

These prior designs have attempted to tackle unfairness in decision making scenarios where the decisions in training data are *biased* (*i.e.*, groups of people with certain sensitive attributes may have historically received unfair treatment) and there is no ground truth about the *correctness* of historical decisions (*i.e.*, one cannot tell whether a decision in training data was right or wrong). However, when the ground truth for historical decisions is available, disproportionately beneficial outcomes for certain sensitive attribute value groups can be justified and explained by the means of the ground truth. Disparate impact would not be a suitable notion of unfairness in such scenarios.

In this paper, we propose an alternative notion of unfairness, ***disparate mistreatment***, especially well-suited for scenarios where ground truth is available for decisions in training data. We identify a decision making process as suffering from disparate mistreatment with respect to a given sensitive attribute (*e.g.*, race) if the *misclassification rates* for groups of people having different values of the sensitive attribute (*e.g.*, blacks and whites) are different. For example, in the case of NYPD Stop-question-and-frisk program (SQF),[1] where people are stopped on the suspicion of having a weapon [20], having different weapon discovery rates for different races would constitute a case of disparate mistreatment.

In addition to *all* misclassifications in general, depending on the application scenario, one might want to define disparate mistreatment with respect to different kinds of misclassifications: false positives, false negatives, or both. For example, in pretrial risk assessments, the decision making process might only be required to ensure that false positive rates are equal since it may be more acceptable to let a guilty person go, rather than incarcerating an innocent person [6, 18]. [2] On the other hand, in loan approval systems, one might instead favor a decision making process in which the false negative rates are equal, to ensure that deserving (positive class) people with a certain sensitive attribute value are not denied (negative class) loans disproportionately. Finally, in the NYPD SQF program, one may argue for a decision making process in which both false positive and false negative rates (on stopping someone on the suspicion of carrying a weapon) are the same for both blacks and whites.

In the remainder of the paper, we first define and formalize disparate treatment, disparate impact and disparate mistreatment in the context of (binary) classification. Then, we introduce an intuitive measure of disparate mistreatment for decision boundary-based classifiers and show that, for a wide variety of linear and nonlinear classifiers, this mea-

[1] https://en.wikipedia.org/wiki/Stop-and-frisk_in_New_York_City

[2] *"It is better that ten guilty persons escape than that one innocent suffer"*—William Blackstone

User Attributes			Ground Truth (Has Weapon)	Classifier's Decision to Stop		
Sensitive	Non-sensitive					
Gender	Clothing Bulge	Prox. Crim.		C_1	C_2	C_3
Male 1	1	1	✓	1	1	1
Male 2	1	0	✓	1	1	0
Male 3	0	1	✗	1	0	1
Female 1	1	1	✓	1	0	1
Female 2	1	0	✗	1	1	1
Female 3	0	0	✓	0	1	0

	Disp. Treat.	Disp. Imp.	Disp. Mist.
C_1	✗	✓	✓
C_2	✓	✗	✓
C_3	✓	✗	✗

Figure 1: Decisions of three fictitious classifiers (C_1, C_2 and C_3) on whether (1) or not (0) to stop a pedestrian on the suspicion of possessing an illegal weapon. Gender is a sensitive attribute, whereas the other two attributes (suspicious bulge in clothing and proximity to a crime scene) are non-sensitive. Ground truth on whether the person is actually in possession of a weapon is also shown.

sure can be incorporated into their formulation as a convex-concave constraint. The resulting formulation can be solved efficiently using recent advances in convex-concave programming [21]. Finally, we experiment with synthetic as well as real world datasets and show that our methodology can be effectively used to avoid disparate mistreatment.

2. BACKGROUND AND RELATED WORK

In this section, we first explain the three different notions of unfairness in decision making scenarios using an illustrative example. Next, we briefly discuss related work on removing two of the three notions namely, *disparate treatment* and *disparate impact*. To the best of our knowledge, there is no prior literature on removing the third notion of unfairness, *disparate mistreatment*, from automated decision making systems.

Disparate mistreatment. Intuitively, disparate mistreatment can arise in any learning scenario, where the outputs (decisions) of the trained algorithm are not perfectly (i.e., 100%) accurate. For instance, consider learning binary classification using logistic regression. If the data items in the two classes are not linearly separable (as is often the case in many real-world application scenarios), the learned algorithm will misclassify (i.e., produce false positives, false negatives, or both) some data items. If the learned algorithm misclassifies people of different genders (males and females) or races (blacks and whites) at different rates, then we refer to such decision making as unfair due to disparate mistreatment.

We illustrate decision making scenarios with and without disparate mistreatment in Figure 1. The scenario involves making a decision to either stop or not stop a pedestrian using their features such as gender, bulge in clothing and proximity to a crime scene. The "ground truth" data for whether a pedestrian actually possesses a weapon is also shown. We show decisions made by three different classifiers $\mathbf{C_1}$, $\mathbf{C_2}$ and $\mathbf{C_3}$. We deem $\mathbf{C_1}$ and $\mathbf{C_2}$ as unfair due to disparate mistreatment because their rate of erroneous decisions for men and women are different: $\mathbf{C_1}$ has different false negative rates for males and females (0.0 and 0.5, respectively), whereas $\mathbf{C_2}$ has different false positive rates (0.0 and 1.0) as well as different false negative rates (0.0 and 0.5) for males and females.

Disparate treatment. In contrast to disparate mistreatment, disparate treatment occurs when the decisions of a classifier for two users who have the same (or similar) values of non-sensitive attributes, but different values of sensitive attributes, are different. For instance, in Figure 1, we deem $\mathbf{C_2}$ and $\mathbf{C_3}$ to be unfair due to disparate treatment since **C2**'s (**C3**'s) decisions for $Male$ 1 and $Female$ 1 ($Male$ 2 and $Female$ 2) are different even though they have the same values of non-sensitive attributes. Here, disparate treatment corresponds to the very intuitive notion of fairness that two otherwise similar users should not be treated differently solely because of their differences in gender.

Disparate impact. Finally, disparate impact is said to occur when the fraction of classifier decisions that benefit (or do not benefit) users belonging to different sensitive attribute groups are different (here, a user benefits from a decision of not being stopped). For instance, in Figure 1, we deem $\mathbf{C_1}$ as unfair due to disparate impact because the fraction of males and females that were stopped are different (1.0 and 0.66, respectively).

Application scenarios for disparate impact vs. disparate mistreatment. Note that unlike in the case of disparate mistreatment, the notion of disparate impact is independent of the "ground truth" information about the decisions, *i.e.*, whether or not the decisions are correct or valid. Thus, the notion of disparate impact is particularly appealing in application scenarios where ground truth information for decisions does not exist and the decision labels in training data are not reliable and cannot be trusted. Unreliability of training labels for automated decision making systems is particularly concerning in scenarios like recruiting or loan approvals, where biased judgments by humans in the past may be used when training classifiers for the future. In such application scenarios, it is hard to distinguish correct and incorrect decisions, making it hard to assess or use disparate mistreatment as a notion of fairness.

However, in application scenarios where ground truth information for decisions can be obtained, disparate impact can be quite misleading as a notion of fairness. That is, in scenarios where validity of decisions can be reliably ascertained, it would be possible to distinguish disproportionality in decision outcomes for sensitive groups that arises from justifiable reasons (*e.g.*, qualification of the users) and disproportionality in outcomes that arises from non-justifiable reasons (i.e., discrimination against sensitive group members). By requiring decision outcomes to be proportional, disparate impact risks introducing reverse-discrimination against qualified candidates. Such practices have previously been deemed unlawful (*Ricci vs. DeStefano, 2009*) by courts. In contrast, when the correctness of decisions can be determined, disparate mistreatment can not only be accurately assessed,

but also avoids reverse-discrimination, making it a more appealing notion of fairness.

Related Work. There have been a number of studies proposing methods for detecting [9, 15, 16, 19] and removing [9, 12, 16, 22, 23] discrimination in the context of disparate treatment and disparate impact. However, as pointed out earlier, these notions might be less meaningful in scenarios where ground truth decisions are available.

A number of recent studies have pointed out racial disparities in both automated [3] as well as human [10, 20] decision making systems related to criminal justice. For example, a work by Goel et al. [20] detects racial disparities in NYPD SQF program, inspired by a notion of unfairness similar to our notion of disparate mistreatment. More specifically, it uses ground truth (stops leading to successful discovery of a weapon on the suspect) to show that blacks were treated unfairly since false positive rates in stops were higher for them than for whites. The study's findings provide further justification for the need of data-driven decision making systems without disparate mistreatment. Our formalization of disparate mistreatment draws inspiration from the analysis of false positive and false negative rates performed by these prior studies [3, 20].

Finally, a recent study [7] analyzing a recidivism risk assessment dataset pointed out that in scenarios where ground truth validation of decisions is available, ensuring disparate impact may lead to different false positive and false negative rates (*i.e.*, disparate mistreatment). This is in line with our intuition that in the presence of ground truth validation, disparate impact might be an unsuitable notion of fairness.

3. FORMALIZING NOTIONS OF FAIRNESS

In a binary classification task, the goal is to learn a mapping $f(\mathbf{x})$ between user feature vectors $\mathbf{x} \in \mathbb{R}^d$ and class labels $y \in \{-1, 1\}$. Learning this mapping is often achieved by finding a decision boundary $\boldsymbol{\theta}^*$ in the feature space that minimizes a certain loss $L(\boldsymbol{\theta})$, *i.e.*, $\boldsymbol{\theta}^* = \text{argmin}_{\boldsymbol{\theta}}\, L(\boldsymbol{\theta})$, computed on a training set $\{(\mathbf{x}_i, y_i)\}_{i=1}^N$. Then, for a given *unseen* feature vector $\mathbf{x}$, the classifier predicts the class label $\hat{y} = f_{\boldsymbol{\theta}^*}(\mathbf{x}) = 1$ if $d_{\boldsymbol{\theta}^*}(\mathbf{x}) \geq 0$ and $\hat{y} = -1$ otherwise, where $d_{\boldsymbol{\theta}^*}(\mathbf{x})$ denotes the signed distance from $\mathbf{x}$ to the decision boundary. Here, we also assume that each user has an associated sensitive feature z. For ease of exposition, we assume z to be binary taking values in $\{0, 1\}$, however, our setup can be easily generalized to categorical as well as multiple sensitive features.

Given the above terminology, we can formally express the absence of disparate treatment, disparate impact and disparate mistreatment as follows:

Existing notion 1: Avoiding disparate treatment. A binary classifier does not suffer from disparate treatment if:

$$P(\hat{y}|\mathbf{x}, z) = P(\hat{y}|\mathbf{x}), \tag{1}$$

i.e., if the probability that the classifier outputs $\hat{y}$ given a feature vector $\mathbf{x}$ does not change after observing the sensitive feature z, there is no disparate treatment.

Existing notion 2: Avoiding disparate impact. A binary classifier does not suffer from disparate impact if:

$$P(\hat{y} = 1|z = 0) = P(\hat{y} = 1|z = 1), \tag{2}$$

i.e., if the probability that a classifier assigns a user to the positive class, $\hat{y} = 1$, is the same for both values of the sensitive feature z, then there is no disparate impact.

New notion 3: Avoiding disparate mistreatment. A binary classifier does not suffer from disparate mistreatment if the misclassification rates—be it all misclassifications combined, false positives or false negatives—for particular groups of people sharing certain sensitive attributes are the same. That is:

$$P(\hat{y} \neq y|z = 0) = P(\hat{y} \neq y|z = 1), \tag{3}$$

$$P(\hat{y} \neq y|z = 0, y = -1) = P(\hat{y} \neq y|z = 1, y = -1), \tag{4}$$

$$P(\hat{y} \neq y|z = 0, y = 1) = P(\hat{y} \neq y|z = 1, y = 1). \tag{5}$$

4. CLASSIFIERS WITHOUT DISPARATE MISTREATMENT

In this section, we describe how to train decision boundary-based classifiers (*e.g.*, logistic regression, SVMs) that do not suffer from disparate mistreatment. These classifiers generally learn the optimal decision boundary by minimizing a convex loss $L(\boldsymbol{\theta})$. The convexity of $L(\boldsymbol{\theta})$ ensures that a global optimum can be found *efficiently*. In order to ensure that the learned boundary is fair—it does not suffer from disparate mistreatment—one could incorporate the appropriate condition from Eqs. 3-5 (based on which kind of misclassifications disparate mistreatment is being defined for) into the classifier formulation. For example:

$$\begin{array}{ll} \text{minimize} & L(\boldsymbol{\theta}) \\ \text{subject to} & P(\hat{y} \neq y|z = 0) - P(\hat{y} \neq y|z = 1) \leq \epsilon, \\ & P(\hat{y} \neq y|z = 0) - P(\hat{y} \neq y|z = 1) \geq -\epsilon, \end{array} \tag{6}$$

where $\epsilon \in \mathbb{R}^+$ and the smaller ϵ is, the more fair the decision boundary would be. The above formulation ensures that the classifier chooses the optimal decision boundary *within* the space of fair boundaries specified by the constraints. However, since the conditions in Eqs. 3-5 may present many saddle points, solving the constrained optimization problem in Eq. (6) seems difficult.

To overcome the above difficulty, we propose a tractable proxy, inspired by the disparate impact proxy proposed by Zafar et al. [22]. In particular, we propose to measure disparate mistreatment using the covariance between the users' sensitive attributes and the signed distance between the feature vectors of misclassified users and classifier decision boundary, *i.e.*:

$$\begin{aligned} \text{Cov}(z, g_{\boldsymbol{\theta}}(y, \mathbf{x})) &= \mathbb{E}[(z - \bar{z})(g_{\boldsymbol{\theta}}(y, \mathbf{x}) - \bar{g}_{\boldsymbol{\theta}}(y, \mathbf{x}))] \\ &\approx \frac{1}{N} \sum_{i=1}^{N} (z_i - \bar{z})\, g_{\boldsymbol{\theta}}(y_i, \mathbf{x}_i), \end{aligned} \tag{7}$$

where the term $\mathbb{E}[(z-\bar{z})]\bar{g}_{\boldsymbol{\theta}}(\mathbf{x})$ cancels out since $\mathbb{E}[(z-\bar{z})] = 0$ and the function $g_{\boldsymbol{\theta}}(y, \mathbf{x})$ is defined as:

$$g_{\boldsymbol{\theta}}(y, \mathbf{x}) = min(0, y d_{\boldsymbol{\theta}}(\mathbf{x})), \tag{8}$$

$$g_{\boldsymbol{\theta}}(y, \mathbf{x}) = min\left(0, \frac{1+y}{2} y d_{\boldsymbol{\theta}}(\mathbf{x})\right), \text{ or} \tag{9}$$

$$g_{\boldsymbol{\theta}}(y, \mathbf{x}) = min\left(0, \frac{1-y}{2} y d_{\boldsymbol{\theta}}(\mathbf{x})\right), \tag{10}$$

which approximates, respectively, the conditions in Eqs. 3-5). Note that, if a decision boundary satisfies Eqs. 3-5, the covariance defined above for that boundary will be close to zero, *i.e.*, $\text{Cov}(z, g_{\boldsymbol{\theta}}(y, \mathbf{x})) \approx 0$. Moreover, in linear models

for classification, such as logistic regression or linear SVMs, the decision boundary is simply the hyperplane defined by $\boldsymbol{\theta}^T\mathbf{x} = 0$, therefore, $d_{\boldsymbol{\theta}}(\mathbf{x}) = \boldsymbol{\theta}^T\mathbf{x}$.

Given the above proxy, one can rewrite Eq. 6 as:

$$\begin{array}{ll} \text{minimize} & L(\boldsymbol{\theta}) \\ \text{subject to} & \frac{1}{N}\sum_{i=1}^{N}(z_i - \bar{z})\, g_{\boldsymbol{\theta}}(y_i, \mathbf{x}_i) \leq \mathbf{c}, \\ & \frac{1}{N}\sum_{i=1}^{N}(z_i - \bar{z})\, g_{\boldsymbol{\theta}}(y_i, \mathbf{x}_i) \geq -\mathbf{c}, \end{array} \tag{11}$$

where the covariance threshold $\mathbf{c} \in \mathbb{R}^+$ controls how *adherent* to disparate mistreatment the decision boundary should be.[3]

Solving the problem efficiently. While the constraints proposed in Eq. 11 can be an effective proxy for fairness, they are non-convex, making it challenging to efficiently solve the optimization problem in Eq. 11. Next, we will convert these constraints into a Disciplined Convex-Concave Program (DCCP), which can be solved efficiently by using recently proposed heuristics [21].

First, consider the constraint described in Eq. 11, *i.e.*,

$$\sum_{i=1}^{N}(z_i - \bar{z})\, g_{\boldsymbol{\theta}}(y_i, \mathbf{x}_i) \sim \mathbf{c}, \tag{12}$$

where $\sim$ may denote '$\geq$' or '$\leq$'. Also we drop the constant number $\frac{1}{N}$ for the sake of simplicity. Since the sensitive feature z is binary, *i.e.*, $z \in \{0, 1\}$, we can split relation (12) into two parts as:

$$\sum_{i=1}^{N_0}(0 - \bar{z})\, g_{\boldsymbol{\theta}}(y_i, \mathbf{x}_i) + \sum_{i=1}^{N_1}(1 - \bar{z})\, g_{\boldsymbol{\theta}}(y_i, \mathbf{x}_i) \sim \mathbf{c}, \tag{13}$$

where N_0 and N_1 are the number of training points taking value $z = 0$ and $z = 1$, respectively. Moreover, since $\bar{z} = \frac{(0\times N_0)+(1\times N_1)}{N} = \frac{N_1}{N}$, we can rewrite 13 as:

$$\frac{-N_1}{N}\sum_{i=1}^{N_0} g_{\boldsymbol{\theta}}(y_i, \mathbf{x}_i) + \frac{N_0}{N}\sum_{i=1}^{N_1} g_{\boldsymbol{\theta}}(y_i, \mathbf{x}_i) \sim \mathbf{c}, \tag{14}$$

which, given that $g_{\boldsymbol{\theta}}(y, \mathbf{x})$ is convex in $\boldsymbol{\theta}$, results into a convex-concave (or, difference of convex) function. Then, substituiting Eq. 14 into Eq. 11, we obtain:

$$\begin{array}{ll} \text{minimize} & L(\boldsymbol{\theta}) \\ \text{subject to} & \frac{-N_1}{N}\sum_{i=1}^{N_0} g_{\boldsymbol{\theta}}(y_i, \mathbf{x}_i) + \frac{N_0}{N}\sum_{i=1}^{N_1} g_{\boldsymbol{\theta}}(y_i, \mathbf{x}_i) \leq \mathbf{c} \\ & \frac{-N_1}{N}\sum_{i=1}^{N_0} g_{\boldsymbol{\theta}}(y_i, \mathbf{x}_i) + \frac{N_0}{N}\sum_{i=1}^{N_1} g_{\boldsymbol{\theta}}(y_i, \mathbf{x}_i) \geq -\mathbf{c}, \end{array} \tag{15}$$

which is a Disciplined Convex-Concave Program (DCCP) for any convex loss $L(\boldsymbol{\theta})$, and can be efficiently solved using well-known heuristics such as the one presented in [21]. Next, we particularize the formulation given by Eq. 15 for a logistic regression classifier [5].[4]

Logistic regression without disparate mistreatment. In logistic regression, the optimal decision boundary $\boldsymbol{\theta}^*$ can be found by solving a maximum likelihood problem of the form $\boldsymbol{\theta}^* = \text{argmin}_{\boldsymbol{\theta}} -\sum_{i=1}^{N}\log p(y_i|\mathbf{x}_i, \boldsymbol{\theta})$ in the training phase. Hence, a fair logistic regressor can be trained by solving the following constrained optimization problem:

$$\begin{array}{ll} \text{minimize} & -\sum_{i=1}^{N}\log p(y_i|\mathbf{x}_i, \boldsymbol{\theta}) \\ \text{subject to} & \frac{-N_1}{N}\sum_{i=1}^{N_0} g_{\boldsymbol{\theta}}(y_i, \mathbf{x}_i) + \frac{N_0}{N}\sum_{i=1}^{N_1} g_{\boldsymbol{\theta}}(y_i, \mathbf{x}_i) \leq \mathbf{c} \\ & \frac{-N_1}{N}\sum_{i=1}^{N_0} g_{\boldsymbol{\theta}}(y_i, \mathbf{x}_i) + \frac{N_0}{N}\sum_{i=1}^{N_1} g_{\boldsymbol{\theta}}(y_i, \mathbf{x}_i) \geq -\mathbf{c}. \end{array} \tag{16}$$

[3]Note that if one wants to have *both* equal false positive and equal false negative rates, one can apply separate constraints with $g_{\boldsymbol{\theta}}(y, \mathbf{x})$ defined in both Eq. 9 and Eq. 10.

[4]Our fairness constraints can be easily incorporated to other boundary-based classifiers such as (non)linear SVMs.

5. EVALUATION

In this section, we conduct experiments on synthetic as well as real world datasets to evaluate the effectiveness of our scheme in controlling disparate mistreatment. To this end, we first generate several *synthetic* datasets to simulate different variations of disparate mistreatment and show that our method can effectively remove disparate mistreatment in each of the variations, often at small cost in terms of accuracy. Then, we conduct experiments on ProPublica COMPAS dataset [14] to show the effectiveness of our method on a *real world* dataset.

In this section, for the sake of simplicity, we assume the sensitive attribute to be binary. However, our methodology and the evaluation can be easily generalized to non-binary, and various sensitive attributes. Additionally, to ensure robustness of experimental findings, for all of the datasets, we repeatedly split the data into training (50%) and testing (50%) folds 5 times and report the average numbers for accuracy and fairness.

Evaluation metrics. Following the definitions of disparate mistreatment provided in Section 3, we define the following metrics to quantify fairness:

$$D_{FP} = P(\hat{y} \neq y|z = 0, y = -1) - P(\hat{y} \neq y|z = 1, y = -1),$$
$$D_{FN} = P(\hat{y} \neq y|z = 0, y = +1) - P(\hat{y} \neq y|z = 1, y = +1),$$

where the closer the values of D_{FP} and D_{FN} to 0, the lower the amount of disparate mistreatment. More specifically, $D_{FP} = 0$ ($D_{FN} = 0$) ensures the fulfillment of Eq. 4 (Eq. 5), *i.e.*, it ensures fairness with respect to the false positive (negative) rate. Additionally, in order to ensure fairness with respect to both false positive and false negative rates, *i.e.*, to satisfy Eq. 3, both D_{FP} and D_{FN} should be zero.

Finally, in the synthetic dataset experiments, we will refer to sensitive attribute values group $z = 0$ as protected group, and $z = 1$ as non-protected group.

Baseline. As pointed out earlier, this work is the first to introduce a mechanism for removing disparate mistreatment in the outcomes of a classification task. Since there are no methods available for performance comparison, we ourselves propose a baseline, which tries to remove disparate mistreatment by introducing different penalties for misclassified data points with different sensitive attribute values. This baseline consists of two main steps. First, we train a (unfair) logistic regressor over the training data and select the set of misclassified data points sharing a sensitive attribute value that presents the higher error rate. For example, if one wants to control for fairness in terms in false positive rate and $D_{FP} > 0$ in the unfair logistic regressor, then we select the set of misclassified data points in training set such that $z = 0$ and $y = -1$. Second, we iteratively re-train the logistic regressor with different (higher) penalties on this set of data points until a certain fairness level is achieved in the training set, in our example, until $D_{FP} \leq \delta$. Algorithm in Figure 2 particularizes this process to ensure fairness in terms of false positive rate. This process can be intuitively

Input: Training set $\{(\mathbf{x}_i, y_i, z_i)\}_{i=1}^N$, $\epsilon > 0$, $\delta > 0$
Output: Fair baseline decision boundary $\boldsymbol{\theta}$
Initialize: Penalty $C = 1$
Train a (unfair) logistic regressor
$\boldsymbol{\theta} = \text{argmin}_{\boldsymbol{\theta}} - \sum_{i=1}^N \log p(y_i|\mathbf{x}_i, \boldsymbol{\theta})$
Compute predicted labels for the training set as
$\hat{y}_i = \text{sign}(d_{\boldsymbol{\theta}}(\mathbf{x}_i))$ and discrimination metric D_{FP}.
Set $d = 0$ if $D_{FP} > 0$, otherwise, $d = 1$.
Select $\mathcal{P} = \{\mathbf{x}_i, y_i, z_i | \hat{y} \neq y_i, z_i = d\}$ and denote $\bar{\mathcal{P}}$ the rest of the training set.
while $D_{FP} > \delta$ **do**
  Increase penalty: $C = C + \epsilon$.
  Train logistic regressor as
  $\boldsymbol{\theta} =$
  $\text{argmin}_{\boldsymbol{\theta}} - \sum_{\bar{\mathcal{P}}} \log p(y_i|\mathbf{x}_i, \boldsymbol{\theta}) - C \sum_{\mathcal{P}} \log p(y_i|\mathbf{x}_i, \boldsymbol{\theta})$
end

Figure 2: Baseline method for removing disparate mistreatment with respect to false positives.

extended to account for fairness in term of false negative rate or for both false positive as well as false negative rates.

5.1 Experiments on synthetic data

In this section, we empirically study the trade-off between fairness and accuracy, in a classifier that suffers from disparate mistreatment in its outcomes. To this end, we first start with a simple scenario in which the classifier is unfair only in terms of either false positive or false negative rate. Then, we focus on a more complex scenario in which the classifier is unfair in terms of both false positive and negative rates.

5.1.1 *Disparate mistreatment on* **only** *false positives or false negatives*

The first scenario accounts for cases in which after training a classifier on the the ground truth labels, we find its outcomes to suffer from disparate mistreatment in terms of only the false positive (or false negative) rate, while they are fair with respect to the other misclassification rate, *i.e.*, $D_{FP} \neq 0$ and $D_{FN} = 0$ (or, alternatively, $D_{FP} = 0$ and $D_{FN} \neq 0$).

Experimental setup. To simulate such a scenario, we first generate 10,000 binary class labels ($y \in \{-1, 1\}$) and corresponding sensitive feature values ($z \in \{0, 1\}$), both uniformly at random, and assign a two-dimensional user feature vector ($\mathbf{x}$) to each of the points. To ensure different distributions for negative class (so that the two sensitive attribute groups have different false positive rates), the user feature vectors are sampled from the following distributions (we sample 2500 points from each distribution):

$$p(\mathbf{x}|z = 0, y = +1) = N([2, 2], [3, 1; 1, 3])$$
$$p(\mathbf{x}|z = 1, y = +1) = N([2, 2], [3, 1; 1, 3])$$
$$p(\mathbf{x}|z = 0, y = -1) = N([1, 1], [3, 3; 1, 3])$$
$$p(\mathbf{x}|z = 1, y = -1) = N([-2, -2], [3, 1; 1, 3])$$

Then, we train a (unconstrained) logistic regression classifier on the data. The classifier is able to achieve an accuracy of 0.85, however, due to difference in feature distributions for the data points with different sensitive attribute values, while $D_{FN} = 0.14 - 0.14 = 0$, $D_{FP} = 0.25 - 0.06 = 0.19$, which constitutes a clear example of disparate mistreatment in terms of false positive rate. Finally, we train several logistic regression classifiers on the same training data subject to fairness constraints on false positive rates, *i.e.*, we train a logistic regressor as in Eq. 16, where $g_{\boldsymbol{\theta}}(y, \mathbf{x})$ is given by Eq. 9. Each classifier constrains the false positive covariance with a multiplicative factor ($\mathbf{c} \in [0, 1]$) of the covariance of the unconstrained classifier. Additionally, we train the proposed baseline as detailed above.

Results. Figure 3 summarizes the results for this scenario by showing (a) the relation between decision-boundary covariance and the false negative rates for both sensitive attribute values; (b) the trade-off between accuracy and fairness; and (c) the decision boundaries for both the unconstrained classifiers (solid) and the fair constrained classifier (dashed). In this figure, we observe that: i) as the fairness constraint value $\mathbf{c}$ goes to zero, the false positive rates for both $z = 0$ and $z = 1$ converge, and hence, the outcomes of the classifier become more fair, *i.e.*, $D_{FP} \to 0$, while D_{FN} remains close to zero (the invariance of D_{FN} may however change depending on the underlying distribution of the data); ii) ensuring lower values of disparate mistreatment leads to a larger drop in accuracy, however, iii) the drop in accuracy for the constrained classifier ($C - D_{FP}$ and $C - D_{FN}$) is still quite modest as compared to the baseline (denoted by $B - D_{FP}$ and $B - D_{FN}$).

5.1.2 *Disparate mistreatment on* **both** *false positives and false negatives*

In this section we consider a more complex scenario, compared to the one presented in previous section, in which the outcomes of the classifier suffer from disparate mistreatment with respect to both false positive as well as false negative rates, *i.e.*, both D_{FP} and D_{FN} are non-zero. This scenario can in turn be split into two cases: (i) D_{FP} and D_{FN} have opposite signs, *i.e.*, the decision boundary favors subjects sharing a sensitive attribute value to be in the positive class (even when such assignments are misclassifications) while assigning the subjects not sharing that value to the negative class and (ii) both D_{FP} and D_{FN} have the same sign, *i.e.*, both error rates are higher for the subjects sharing a sensitive attribute value.

Case I. The first case, where D_{FP} and D_{FN} have opposite signs, occurs when the decision boundary favors one sensitive attribute value group by misclassifying a relatively large number of its negative points into positive class (resulting in a relatively high false positive rate) and disfavors the other group by misclassifying a large number of its positive points into negative class (resulting in a high false negative rate).

To simulate this scenario, we first generate $2,500$ samples from each of the following distributions:

$$p(\mathbf{x}|z = 0, y = 1) = N([2, 0], [5, 1; 1, 5])$$
$$p(\mathbf{x}|z = 1, y = 1) = N([2, 3], [5, 1; 1, 5])$$
$$p(\mathbf{x}|z = 0, y = -1) = N([-1, -3], [5, 1; 1, 5])$$
$$p(\mathbf{x}|z = 1, y = -1) = N([-1, 0], [5, 1; 1, 5])$$

An unconstrained logistic regression classifier on this dataset attains an overall accuracy of 0.78 but leads to a false positive rate of 0.14 and 0.30 (*i.e.*, $D_{FP} = 0.14 - 0.30 = -0.16$) for the protected and the non-protected group; and false negative rates of 0.31 and 0.12 (*i.e.*, $D_{FN} = 0.31 - 0.12 = 0.19$) for the protected and the non-protected group. Finally, we

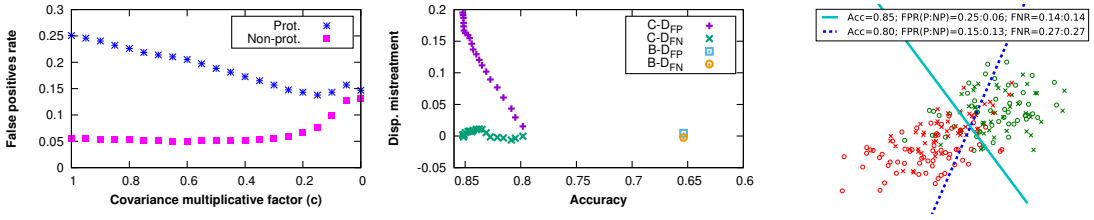

(a) Cov. vs FPR (b) Fairness vs. Acc. (c) Boundaries

Figure 3: [Synthetic data] Panel (a) shows that decreasing the covariance threshold causes the false positive rates for protected and non-protected groups to becomes similar. Panel (b) shows that an increasing amount of fairness corresponds to a steady decrease in accuracy. However, the drop in accuracy for our method ($C - D_{FP}$ and $C - D_{FN}$) is modest as compared to the one achieved by the baseline (denoted by $B - D_{FP}$ and $B - D_{FN}$) for the same amount of fairness. Panel (c) shows the original decision boundary (solid line) and fair decision boundary (dashed line), along with corresponding accuracy and false positive rates for protected (P, crosses) and non-protected (NP, circles) groups. Fairness constraints cause the original decision boundary to rotate such that previously misclassified protected examples are moved into the negative class (decreasing false positives), while well-classified non-protected examples are moved into the positive class (increasing false positives), leading to equal false positive rates for both groups.

train three different type of fair classifiers, with fairness constraints on (i) false positive rates—$g_{\boldsymbol{\theta}}(y, \mathbf{x})$ given by Eq. 9, (ii) false negative rates—$g_{\boldsymbol{\theta}}(y, \mathbf{x})$ given by Eq. 10 and (iii) on both false positive and false negative rates—separate constraints for $g_{\boldsymbol{\theta}}(y, \mathbf{x})$ given by Eq. 9 and Eq. 10.

Results. Figure 4 summarizes the results for this scenario by showing the (top) trade-off between fairness and accuracy for the three constrained classifiers described above, and (bottom) the decision boundaries for the unconstrained classifier (solid) and the constrained fair classifiers. Here, we can observe several interesting patterns. First, removing disparate mistreatment on only false positives causes a rotation in the decision boundary to move *previously misclassified* non-protected examples into negative class, *decreasing* false positive rate. However, in the process, it also moves previously *well-classified* non-protected examples into negative class, *increasing* the false negative rate for the non-protected group. As a consequence, controlling disparate mistreatment on false positives (Figure 4(a)), also removes disparate mistreatment on false negatives. Note that, as similar effect occurs when we control disparate mistreatment only with respect to the false positive rate (Figure 4(b)), and therefore, provides similar results that the constrained classifier for both false positive and negative rates (Figure 4(c)). This effect is explained by the distribution of the data, where the centroids of the clusters for the protected group are shifted with respect to the ones for the non-protected group. Finally, we remark that while the constrained classifiers (denoted by $C - D_{FP}$ and $C - D_{FN}$) remove disparate mistreatment with respect to both false positive and negative rates at a moderate loss in accuracy, the baseline (denoted by $B - D_{FP}$ and $B - D_{FN}$) do not remove disparate mistreatment with respect to both false positive and negative rates in any of the cases and result in a equal or higher loss in accuracy for the same fairness level (*i.e.*, D_{FP} and D_{FN} values).

Case II. The second case account for scenarios in which after training a (unconstrained) classifier, its outcomes results in disparate mistreatment where both D_{FP} and D_{FN} share the same sign, *i.e.*, both error rates are higher for the subjects sharing a certain sensitive attribute value. This scenario is likely to occur in datasets where the subjects with one sensitive attribute value (*e.g.*, females) are harder to classify than the subjects not sharing that sensitive attribute value (*e.g.*, males)—in a extreme case, the two classes may be linearly separable for males, but not for the females.

Experimental setup. To simulate this scenario, we first generate $2,500$ samples from each of the following distribution:

$$p(\mathbf{x}|z = 0, y = 1) = N([1, 2], [5, 2; 2, 5])$$
$$p(\mathbf{x}|z = 1, y = 1) = N([2, 3], [10, 1; 1, 4])$$
$$p(\mathbf{x}|z = 0, y = -1) = N([0, -1], [7, 1; 1, 7])$$
$$p(\mathbf{x}|z = 1, y = -1) = N([-5, 0], [5, 1; 1, 5])$$

Then, we train an unconstrained logistic regression classifier on this dataset. It attains an accuracy of 0.80 but leads to $D_{FP} = 0.33 - 0.08 = 0.25$ and $D_{FN} = 0.26 - 0.11 = 0.15$, which constitutes a clear case of disparate mistreatment in terms of both false positive and negative rates. Then, similar to the previous scenario, we train three different kind of constrained classifiers to remove disparate mistreatment on (i) false positives, (ii) false negatives. and (iii) both.

Results. Figure 5 summarizes the results by showing the (top) the trade-off between accuracy and fairness, and (bottom) the decision boundaries for both the unconstrained classifiers (solid) and the fair constrained classifier (dashed) when accounting for disparate mistreatment with respect to false positive rate, false negative rate or both. In this figure, we observe several interesting patterns. First, controlling disparate mistreatment for only false positives (false negatives), leads to a minor drop in accuracy, but can exacerbate the disparate mistreatment on false negatives (false positives). For example, while the decision boundary is moved to control for false negative disparate mistreatment to ensure that more protected examples are well-classified in positive class (reducing false negative rate), it also moves previously well-classified negative examples into the positive class, hence increasing the false positive rate. A similar phenomenon occur when controlling disparate mistreatment with respect to only false positive rate. As a consequence, controlling for both types of disparate mistreatment simultaneously brings D_{FP} and D_{FN} close to zero, but causes a large drop in accuracy. Second, while the constrained classifiers always remove the targeted disparate mistreatment often at a low cost in accuracy, the baseline is not able to

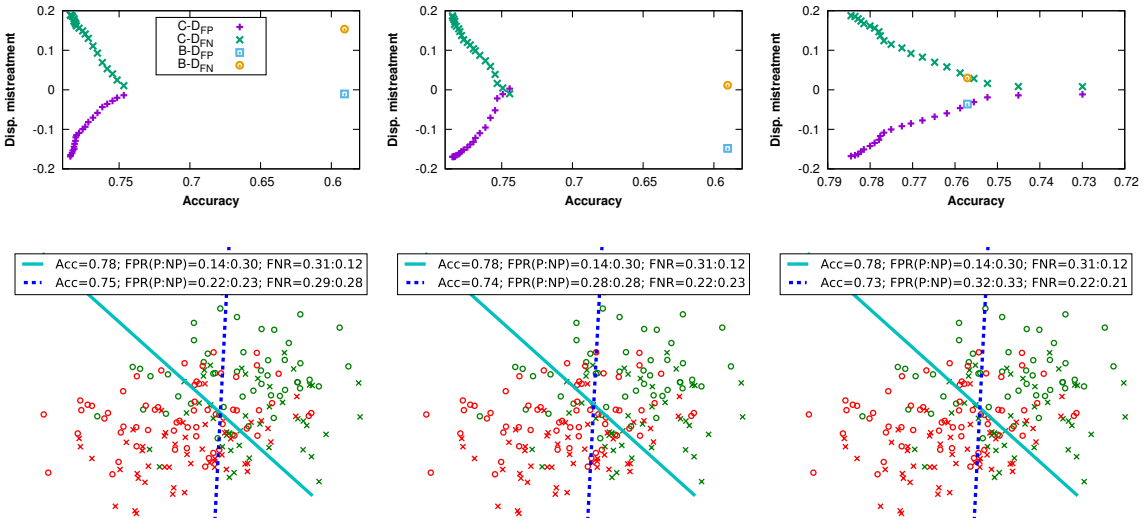

(a) FPR Constraints (b) FNR Constraints (c) Both Constraints

Figure 4: [Synthetic dataset] D_{FP} and D_{FN} have opposite signs. Removing disparate mistreatment on false positives also helps remove disparate mistreatment on false negatives. Removing disparate mistreatment on both at the same time (Panel (c)) leads to very similar results.

Race	Yes	No	Total
Black	1,661(52%)	1,514(42%)	3,175(100%)
White	8,22(39%)	1,281(61%)	2,103(100%)
Total	2,483(47%)	2,795(53%)	5,278(100%)

Table 1: Statistics on COMPAS data showing whether an offender recidivated within two years or not.

Feature	Description
Age Category	< 25, between 25 and 45, > 45
Gender	Male or Female
Race	White or Black
Priors Count	0–37
Charge Degree	Misconduct or Felony
2-year-rec. (target feature)	Whether (+ve) or not (-ve) the defendant recidivated within two years

Table 2: Description of features used from ProPublica COMPAS data.

ensure fairness in the final setting (controlling for both) and also presents a larger drop in accuracy in all settings.

5.2 Real world dataset: ProPublica COMPAS

In this section, we experiment with COMPAS risk assessment dataset compiled by ProPublica [14] and show that our method can significantly reduce disparate mistreatment at a modest cost in terms of accuracy.

Dataset and experimental setup. ProPublica compiled a list of all offenders screened through Correctional Offender Management Profiling for Alternative Sanctions (COMPAS) tool[5] in Broward County, Florida during 2013-2014. The data includes information about offenders' demographic features (gender, race, age), criminal history (charge for which the person was arrested, number of prior offenses) and the risk score assigned to the offender by COMPAS. ProPublica also collected the *ground truth* on whether or not these individuals actually recidivated within two years after screening. For more information about how the data was collected, we point the reader to [2]. Finally, for the sake of simplicity, in this analysis, we will only consider a subset of offenders whose race was either black or white. The information about recidivism rates for each race are included in Table 1.

Using this ground truth, we build a logistic regression classifier to predict whether (positive class) or not (negative class) an offender would recidivate within two years. The set of features used in the classification task are described in Table 2, which is the same as the features used by ProPublica for their analysis of modeling risk scores. [6] The logistic regression classifier leads to an accuracy of 0.668 and AUC of 0.721. However, the classifier yields in a false positive rate of 0.34 and 0.15 for respectively blacks and whites (*i.e.*, $D_{FP} = 0.19$), and false negative rates for respectively blacks and whites of 0.32 and 0.55 (*i.e.*, $D_{FN} = -0.23$). This results constitute a clear case of disparate mistreatmentin terms of both false positive and negative rates, where one sensitive feature group (blacks) is put on relative disadvantage by the classifier by disproportionally misclassifying negative (good, did not recidivate) examples into positive (bad) class (false positives), while the other group (whites) is put on relative advantage by disproportionally misclassify-

[5] COMPAS tries to predict the risk of recidivism for a criminal offender. The process consists of analyzing answers to 137 questions pertaining to the offender's criminal history and behavioral patterns [1]. The output of the system is a decile score (1–10) predicting the risk of recidivism.

[6] Notice that goal of this section is not to analyze the best sets of features for recidivism prediction, rather, we focus on showing that our method can effectively remove disparate mistreatment in a given dataset. Hence, we chose to use the same set of features as used by ProPublica for their analysis.

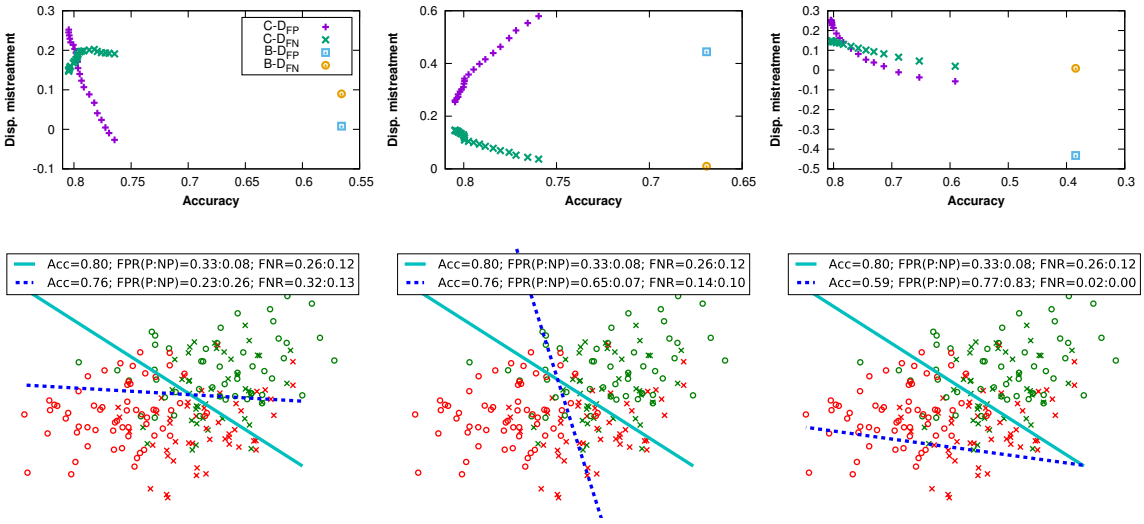

(a) FPR Constraints (b) FNR Constraints (c) Both Constraints

Figure 5: [Synthetic dataset] D_{FP} and D_{FN} have same signs. Removing disparate mistreatment on false positives can potentially increase disparate mistreatment on false negatives. Removing both types of disparate mistreatment at the same time (in (c)) causes a larger drop in accuracy.

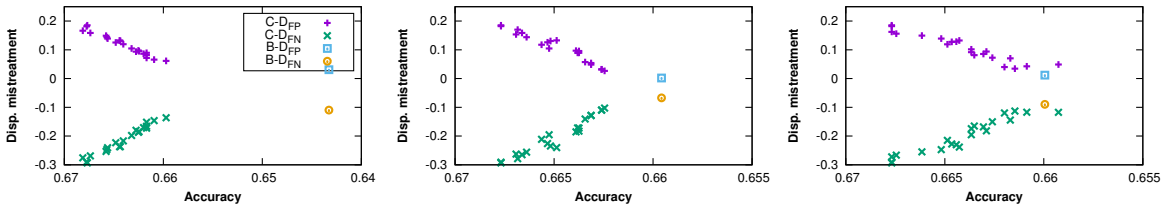

(a) FPR Constraints (b) FNR Constraints (c) Both Constraints

Figure 6: Removing disparate mistreatment on false positives / false negatives / both in ProPublica COMPAS data.

ing positive (bad, recidivated) examples into negative (good) class (false negatives). Note therefore that this scenario resembles our synthetic example Case I in Section 5.1.2.

Finally, we train logistic regression classifiers with three types constraints: false positives, false negatives, and both.

Results. Figure 6 summarizes the results by showing the trade-off between fairness and accuracy for the three constrained classifiers and the proposed baselines. Similarly to the results in Section 5.1.2, we observe here that the logistic regressor with constraints on false positive (or alternatively negative) rate is also controlling disparate mistreatmentin terms of false negative (or alternatively positive) rate, and it does so at a lower drop in accuracy than the baseline. However, when we control disparate mistreatmentin terms of both false positive and negative rates, the baseline provides a slightly better trade-off between accuracy and fairness than the proposed constrained logistic regressor. Additionally, we observe that the constrained classifiers (as well as the baseline classifier) do not remove completely disparate mistreatment, *i.e.*, they don't achieve zero D_{FP} or/and D_{FN} in any of the cases. This is probably due to the relatively small size of the dataset (2,639 examples in the training set), which hinders a robust estimate of misclassification covariance (Eqs. 9 and 10). This highlights the fact that our method can suffer from reduced performance for small datasets. In scenarios with sufficiently large training datasets, we expect more reliable estimates of covariance, and hence, a better performance from our method.

6. DISCUSSION AND FUTURE WORK

As shown in Section 5, our method provides a flexible trade-off between disparate mistreatment-based fairness and accuracy. However, we would like to point out that the current formulation of fairness constraints may suffer from the following two main limitations. Addressing these two limitations would be an interesting avenue for future work.

First, the proposed formulation to train fair classifiers is not a convex program, but a disciplined convex-concave program (DCCP). As a consequence, our convex-concave formulation can be efficiently solved using recently proposed heuristics [21], which work well in practice. However, unlike convex optimization, these heuristic-based methods do not provide any guarantees about the global optimality of the solution. Secondly, since computing the analytical covariance in fairness constraints is not a trivial task, we approximate the covariance using Monte-Carlo approximation on training

set (Eq. 7). While this approximation is expected to work well when a reasonable amount of training data is provided, it might be inaccurate for smaller datasets. This problem is exacerbated, in comparison to [22], by the fact that we are not using the whole training set to approximate the covariance but only the misclassified points. Therefore, for small error rates, our scheme may not be able to provide good estimates of covariance and lead to sub-optimal solutions.

7. REFERENCES

[1] https://www.documentcloud.org/documents/2702103-Sample-Risk-Assessment-COMPAS-CORE.html.

[2] https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm.

[3] J. Angwin, J. Larson, S. Mattu, and L. Kirchner. Machine bias: There's software used across the country to predict future criminals. and it's biased against blacks. https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing, 2016.

[4] S. Barocas and A. D. Selbst. Big Data's Disparate Impact. *California Law Review*, 2016.

[5] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.

[6] W. Blackstone. *Commentaries on the Laws of England*. 1760.

[7] A. Chouldechova. Fair Prediction with Disparate Impact:A Study of Bias in Recidivism Prediction Instruments. *arXiv:1610.07524*, 2016.

[8] EFF. http://tinyurl.com/EFF-Web, 2005.

[9] M. Feldman, S. A. Friedler, J. Moeller, C. Scheidegger, and S. Venkatasubramanian. Certifying and removing disparate impact. In *KDD*, 2015.

[10] J. M. Greg Ridgeway. Doubly robust internal benchmarking and false discovery rates for detecting racial bias in police stops. *Journal of the American Statistical Association*, 2009.

[11] M. Hardt. How Big Data is Unfair: Understanding Sources of Unfairness in Data Driven Decision Making. *Medium*, 2014.

[12] F. Kamiran and T. Calders. Classification with No Discrimination by Preferential Sampling. In *BENELEARN*, 2010.

[13] T. Kamishima, S. Akaho, H. Asoh, and J. Sakuma. Fairness-aware Classifier with Prejudice Remover Regularizer. In *PADM*, 2011.

[14] J. Larson, S. Mattu, L. Kirchner, and J. Angwin. https://github.com/propublica/compas-analysis, 2016.

[15] B. T. Luong, S. Ruggieri, and F. Turini. kNN as an Implementation of Situation Testing for Discrimination Discovery and Prevention. In *KDD*, 2011.

[16] D. Pedreschi, S. Ruggieri, and F. Turini. Discrimination-aware Data Mining. In *KDD*, 2008.

[17] J. Podesta, P. Pritzker, E. Moniz, J. Holdren, and J. Zients. Big Data: Seizing Opportunities, Preserving Values. *Executive Office of the President. The White House.*, 2014.

[18] Reasonable Doubt. wikipedia.org/wiki/Reasonable_doubt.

[19] A. Romei and S. Ruggieri. A Multidisciplinary Survey on Discrimination Analysis. *KER*, 2014.

[20] R. S. Sharad Goel, Justin M. Rao. Precinct or Prejudice? Understanding Racial Disparities in New York City's Stop-and-Frisk Policy. *Annals of Applied Statistics*, 2015.

[21] X. Shen, S. Diamond, Y. Gu, and S. Boyd. Disciplined Convex-Concave Programming. *arXiv:1604.02639*, 2016.

[22] M. B. Zafar, I. V. Martinez, M. G. Rodriguez, and K. P. Gummadi. Learning Fair Classifiers. *arXiv:1507.05259*, 2016.

[23] R. Zemel, Y. Wu, K. Swersky, T. Pitassi, and C. Dwork. Learning Fair Representations. In *ICML*, 2013.